

The Painter's Problem: Covering a Grid with Colored Connected Polygons

Arthur van Goethem¹, Irina Kostitsyna¹, Marc van Kreveld²,
Wouter Meulemans¹, Max Sondag¹, and Jules Wulms¹✉

¹ Department of Mathematics and Computer Science, TU Eindhoven,
Eindhoven, The Netherlands

{a.i.v.goethem,i.kostitsyna,w.meulemans,m.f.m.sondag,
j.j.h.m.wulms}@tue.nl

² Department of Information and Computing Sciences, Utrecht University,
Utrecht, The Netherlands

m.j.vankreveld@uu.nl

Abstract. Motivated by a new way of visualizing hypergraphs, we study the following problem. Consider a rectangular grid and a set of colors χ . Each cell s in the grid is assigned a subset of colors $\chi_s \subseteq \chi$ and should be partitioned such that for each color $c \in \chi_s$ at least one piece in the cell is identified with c . Cells assigned the empty color set remain white. We focus on the case where $\chi = \{\text{red, blue}\}$. Is it possible to partition each cell in the grid such that the unions of the resulting red and blue pieces form two connected polygons? We analyze the combinatorial properties and derive a necessary and sufficient condition for such a *painting*. We show that if a painting exists, there exists a painting with bounded complexity per cell. This painting has at most five colored pieces per cell if the grid contains white cells, and at most two colored pieces per cell if it does not.

1 Introduction

Hypergraphs are a powerful structure to represent unordered set systems. In general, there are a number of elements (vertices of the hypergraph) and a number of different subsets over these elements (the hyperedges of the graph). The purpose of visualizing hypergraphs is to clarify the various set relations between the hyperedges. There are, roughly speaking, two strands of hypergraph visualizations: those where the position of the elements is fixed (e.g. [2, 7, 8, 15]), and those where the positions can be chosen by the layout algorithm (e.g. [10, 18, 19]). For a more detailed overview and in-depth classification of set visualization methods we refer to the survey by Alsallakh et al. [4]. Though some methods aim to

This work was initiated at the 2nd Workshop on Applied Geometric Algorithms (AGA 2017) supported by the Netherlands Organisation for Scientific Research (NWO), 639.023.208. AvG is supported by NWO 612.001.102; IK by FRS-FNRS; MvK by NWO 612.001.651; WM and JW by NLeSC 027.015.G02; MS by NWO 639.023.208.

overcome layout complexity by replicating elements (e.g. [3,10]), we focus on a visualization using a single representation for each element.

In theoretic research on drawing hypergraphs (e.g. [6,12]), the (often implicit) assumption is that the representations of two sets may cross at common vertices. Such crossings are not deemed problematic as most visual encodings rely on the local *nesting* of intersecting polygons (in line with the prototypical Venn and Euler diagrams [5] and similar visual overlays [7,8,15]) to identify set memberships. Nesting, however, gives a strong visual cue of containment and may result in misleading visual representations implying containment relationships between hyperedges. A rendering style without nesting is one suggested for Kelp Diagrams [8]. However, its cluttered appearance caused it not to feature in the later extension, KelpFusion [15].

One of the most well-established quality criteria of graph drawings is planarity (see e.g. [16,17]). When nested encodings are used, a planar drawing relates to finding a planar support [6]: a planar (regular) graph such that the vertices of each hyperedge induce a connected subgraph in the support. Deciding whether a planar support exists is possible for some simple support classes (see [6] for a discussion), but is already NP-hard for 2-outerplanar support graphs [6]. Optimizing hypergraph supports for total graph length without planarity constraints is NP-hard, but approximation algorithms exist [1,11].

Representations that do not require nesting are edge-based drawings [13] or the equivalent Zykov representation [22], for which notions of planarity follow readily from the standard notion for regular graphs.

Instead we suggest a visual design that uses *disjoint polygons* to present hyperedges: vertices are represented as simple geometric primitives (e.g. a square or circle); hyperedges are represented as connected polygons that overlaps only and all its incident vertices; and all such polygons are pairwise disjoint. As illustrated in Fig. 1, our disjoint-polygons encoding is stronger as it can visualize some hypergraphs that are not Zykov-planar, whereas any Zykov-planar hypergraph admits a disjoint-polygons representation. We can use vertices to “pass in between” the representations of other hyperedges, though not as flexibly as is allowed for planar supports: the polygons must remain disjoint.




Fig. 1. A hypergraph that is not Zykov-planar (top) but has a disjoint-polygons drawing (bottom).

Contributions. We investigate the properties of drawing hypergraphs using disjoint polygons. Motivated by moving towards a set visualization in a geographic small multiples or grid map (see e.g. [14,23]), we specifically study the variant where each element has a fixed location, being a cell in a rectangular grid. As an initial exploration we focus on the 2-color case, where each cell is either red, blue, both (purple), or uncolored (white). We thus aim to partition each purple cell into red and blue pieces, such that the resulting pieces of a single color form a connected polygon. We derive a necessary and sufficient condition to efficiently

recognize whether an instance is solvable. For solvable instances, we bound the number of colored pieces within each cell by a small constant and show that these bounds are tight. Due to space constraints, some proofs have been shortened or omitted; for full proofs, please refer to the ArXiv version [21].

2 Preliminaries

We define a *k-colored grid* Γ as a rectangular grid, in which each cell s has a set of associated colors $\chi_s \subseteq \{1, \dots, k\}$. A *fully k-colored grid* is the case where $\chi_s \neq \emptyset$ for all cells s . Throughout this paper, we primarily investigate 2-colored grids and use *colored grid* to refer to the 2-colored case, unless indicated otherwise. We refer to the two colors as *(r)ed* and *(b)lue*; cells for which $\chi_s = \{r, b\}$ are called *(p)urple*. Cells with no associated colors are *white*.

A region is a maximal set of cells that have the same color assignment (*r*, *b*, or *p*) and where every cell s in the region is connected via adjacent cells to every other cell s' in the region. Cells are considered adjacent if they are horizontally or vertically adjacent.

A *panel* π_s for cell s (with $\chi_s \neq \emptyset$) maps each color $c \in \chi_s$ to a (possibly disconnected) area $\pi_s(c)$ such that these partition the cell: that is, $\bigcup_{c \in \chi_s} \pi_s(c) = s$ and $\pi_s(c_1) \cap \pi_s(c_2) = \emptyset$ for colors $c_1 \neq c_2$. A *painting* Π of a *k-colored grid* consists of panels π_s for each cell s with $\pi_s(c) \neq \emptyset$ for each $c \in \chi_s$ and $\pi_s(c) = \emptyset$ otherwise. We call a painting *connected* if each color forms a connected polygon: that is, $\bigcup_{s \in \Gamma} \pi_s(c)$ is a connected polygon for each color $c \in \{1, \dots, k\}$. For this definition, two cells sharing only a corner are *not* considered connected. Our primary interest is in connected paintings: in the remainder, we use painting to indicate a connected painting.

3 Characterizing Colored Grids with a Painting

In this section we show how to test whether a 2-colored grid admits a painting and how to find a painting if one exists. As all completely red, blue, and white panels are fixed, finding a painting reduces to finding partitions of purple cells that ensure that the resulting red and blue polygon are connected. We show that this connectivity is of key importance: if we can find suitable connections though the purple regions, then we can also create a partition that results in a valid panel for each cell in the purple regions.

We capture the connectivity options for the red and blue polygon using two embedded graphs, G_r and G_b . We construct these graphs in three steps:

1. Connect red (blue) regions that are adjacent along a purple region's boundary.
2. Remove holes from the purple regions by inserting connections (Sect. 3.3).
3. Construct G_r and G_b using a gadget for purple regions (Sects. 3.1 and 3.2).

For the first step, observe that consecutive (not necessarily distinct) adjacent regions of the same color can always be safely connected via the purple region's




Fig. 2. Safe connections between adjacent same-color regions. (Color figure online)





Fig. 3. 2-colored grid with 4 regions around each purple region and corresponding graphs G_r and G_b . (Color figure online)



boundary without restricting the connectivity options for the other color (see Fig. 2). After the first two steps, we represent the remaining red (blue) regions as vertices in G_r (G_b). Edges in G_r and G_b represent connection options through purple regions; intersections indicate a choice to connect either blue or red regions through (part of) a purple region. The gadget for purple regions with many adjacent red and blue regions also requires some additional vertices in these graphs. We prove that these graphs admit a simple characterization of 2-colored grids that admit a painting, as captured in the theorem below.

Theorem 1. *A 2-colored grid Γ admits a painting if and only if the corresponding graphs G_r and G_b are each other’s exact duals: there is exactly one blue vertex in every red face and there is exactly one red vertex in every blue face.*

For explanatory reasons we start with the simplest case: purple regions with at most four neighbors and without holes (Sect. 3.1). Subsequently, we alleviate the assumption on the number of neighbors (Sect. 3.2) and permit holes in the purple regions, by showing how to perform Step 2 (Sect. 3.3).

3.1 Simple Purple Regions

We assume that Step 1 has been performed and a purple region has no holes and at most four adjacent regions. The adjacent red and blue regions of a purple region P form an ordered cyclic list as they appear along the boundary of P and alternate in color (due to Step 1). Let $\kappa(P)$ denote the length of this list for P . κ is even due to color alternation, and by assumption here $\kappa(P) \leq 4$. There can be duplicates in this list as the same red or blue region can touch P multiple times.

Every purple region with $\kappa(P) = 2$ can be painted by creating a spanning tree on the centers of the panels of P in one color and connecting it to the corresponding region. The rest of the panels is colored in the other color. We assume these are handled; what remains is to deal with the regions with $\kappa(P) = 4$.

For a purple region P with $\kappa(P) = 4$, we create a red edge in G_r and a blue edge in G_b that intersect: the red edge connects the red vertices corresponding to

the adjacent red regions; the blue edge connects the corresponding blue vertices. There may be multiple edges between two vertices (see Fig. 3). If the same red or blue region touches the purple region twice, the edge is a self-loop. Every red or blue edge intersects exactly one blue or red edge respectively, and G_r and G_b are plane by construction. Using the following lemma we prove the exact characterization of graphs G_r and G_b of a 2-colored grid Γ that admits a painting.

Lemma 1 ([9, 20]). *Let G be a plane graph, G^* its dual and T a spanning tree of G . Then $T^* = \{e^* \mid e \notin T\}$ is a spanning tree of G^* .*

Lemma 2. *A 2-colored grid Γ in which each purple region P has no holes and $\kappa(P) \leq 4$, admits a painting if and only if the corresponding graphs G_r and G_b are each other's exact duals.*

Proof (sketch). We prove that if Γ admits a painting then graphs G_r and G_b are each other's duals using a counting argument. We count the number of edges needed to connect all red and blue regions, and use Euler's formula to show the number of red faces must be equal to the number of blue vertices, and vice versa. The other direction follows from Lemma 1. Having two dual spanning trees (e.g., Fig. 3), simply draw the two spanning trees and for any cell not yet having a blue (red) piece add a crossing-free connection to the blue (red) polygon. \square

3.2 Spiderweb Gadgets

Let us now extend the result in the previous section, by showing how to include purple regions with more than four adjacent regions. For every purple region P with $\kappa(P) > 4$ we construct a *spiderweb* gadget and insert it into the graphs G_r and G_b , such that an argument similar to Lemma 2 can be applied.

A spiderweb gadget W of P with $\kappa(P)/2 = k$ red and k blue alternating adjacent regions consists of $\lfloor k/2 \rfloor + 1$ levels, labeled 0 (outermost) to $\lfloor k/2 \rfloor$ (innermost), see Fig. 4. Each level, except 0 and $\lfloor k/2 \rfloor$, is a cycle of k vertices. The level 0 has k (blue) vertices without any edges between them, and the innermost level $\lfloor k/2 \rfloor$ consists of only a single vertex. The vertices of even levels are blue and labeled with even numbers from 0 to $2k - 2$ clockwise. The vertices of odd levels are red and labeled with odd numbers 1 to $2k - 1$ clockwise.

Each vertex of level ℓ with $2 \leq \ell < \lfloor k/2 \rfloor$ is connected to the vertex with the same label on level $\ell - 2$. The single vertex of level $\lfloor k/2 \rfloor$ is connected to all the vertices of level $\lfloor k/2 \rfloor - 2$. This gives us $2k$ paths starting from levels 0 and 1 to the two innermost levels. We call these paths *spokes*, and refer to them by the label of the corresponding vertices. We embed the two resulting connected components in such a way that they are each other's dual, by making sure that we get a proper clockwise numbering on the vertices of the two outermost levels (see Fig. 4). The vertices on levels 0 and 1 represent respectively the blue and red regions around the purple region P and respect the adjacency order around P .

If a blue (or red) region touches P multiple times, then the corresponding vertices on level 0 (or 1) map to the same region and are in fact one and the same




Fig. 4. Spiderweb gadget for $k = 6$: three blue levels with indices 0, 2, 4, and two red levels with indices 1, 3. (Color figure online)




Fig. 5. Topology of the connections in a purple region and the corresponding bridging paths through a spiderweb gadget. (Color figure online)

vertex in G_b (or G_r). All edges connected to this vertex are consistent with the topology of the nested neighboring regions of P ; they intersect the same edges as they would when they were represented by multiple vertices.

To prove that all possible connections in P , which can occur in a painting Π , can be replicated in a spiderweb gadget W , we define *bridging paths*: let u and v be two vertices on level 0 in W that represent two blue regions that are connected by a painting Π through P . Assume that the clockwise distance from u to v is not greater than k , that is, if u has label x then v has label $(x + 2i) \bmod 2k$ for some $1 \leq i \leq \lfloor k/2 \rfloor$. To connect u and v with a bridging path, we start from u , go to level $2\lfloor(i + 1)/2\rfloor$ along the spoke x , take a shortest path within the level $2\lfloor(i + 1)/2\rfloor$ from the vertex with label x to the vertex with label $(x + 2i) \bmod 2k$, and move along the spoke $(x + 2i) \bmod 2k$ to vertex v . If there are two possible shortest paths, we take the clockwise path.

The same kind of path can be constructed for a pair of red vertices, but starting from level 1, going to level $2\lfloor i/2 \rfloor + 1$, and moving back to level 1. We now show that connecting different blue and red regions using bridging paths within the spiderweb gadgets results in blue trees and red trees, such that no pair of a blue and a red edge intersect (see Fig. 5 for an example).

By performing a case analysis on the possible red and blue pairs of adjacent regions to be connected, we can prove that the following lemma holds.

Lemma 3. *Consider a painting Π in which two blue and two red regions, adjacent to a purple region P , are connected through P . The corresponding vertices in the spiderweb gadget W of P can be connected by non-intersecting bridging paths.*

With spiderweb gadgets and the above lemma, we now strengthen Lemma 2 to the following lemma, without a condition on κ , and prove it in a similar way.

Lemma 4. *A 2-colored grid Γ in which each purple region has no holes admits a painting if and only if the corresponding G_r and G_b are each other’s exact duals.*




Fig. 6. An annulus-type purple region with adjacent blue and red regions, both inside and outside. (Color figure online)




Fig. 7. By adding edges (v_{in}, y_{in}) and (u_{out}, x_{out}) we reconnect the disconnected subpolygons formed by removing cross-annulus connections γ_v and γ_x . (Color figure online)

3.3 Purple Regions with Holes

We may also have purple regions with holes (see Fig. 6). We show that the number of holes can be reduced without affecting the solvability. For simplicity of explanation we assume a region with a single hole (an annulus); regions with more holes can be reduced by considering only connections to the outer boundary.

Let P be a purple annulus. Any painting subdivides P into a number of colored simple components. Each component of color c connects one or more regions of color c on the boundary of P . The existence of a painting is defined only by the connectivity structure of these components. The connectivity of a component can be represented (transitively) using a set of non-intersecting simple paths (*connections*) each connecting two regions on the boundary. Let a *cross-annulus connection* γ_x be a connection between a region x_{in} on the inside of the annulus and a region x_{out} on the outside of the annulus. A (connectivity) *structure* is a maximal set of (pairwise non-intersecting) connections in P that can be extended to a valid painting. Let C_S be the set of cross-annulus connections in a given structure S . We assume the annulus is not degenerate, so red and blue regions exist both inside and outside the annulus.

Lemma 5. *If a structure S exists with two adjacent cross-annulus connections γ_x and γ_y of the same color, possibly separated by non-crossing connections, then there also exists a structure S' where $C_{S'} = C_S \setminus \{\gamma_y\}$.*

Lemma 6. *If there exists a structure S with $|C_S| > 3$ and all cross-annulus connections are alternating in color, then there also exists a structure S' with $|C_{S'}| - 2$ cross-annulus connections.*

Proof. Let $\gamma_u, \gamma_v, \gamma_x,$ and γ_y be four consecutive cross-annulus connections. W.l.o.g., assume γ_u and γ_x are red and γ_v and γ_y are blue. We remove γ_v and γ_x from the structure separating both the red and blue into two components. For both colors one component is still connected to the remaining cross-annulus connection γ_u , respectively γ_y . The disconnected components cannot both be




Fig. 8. (a) Initial configuration with several connections covering r_{out} . (b) Rerouting the blue connections, introducing γ_z , and rerouting the intersecting red connection leaves only one intersecting (blue) connection. (c) As the blue disconnected component cannot be covered by the new red connection, we can always connect it back to γ_x . (Color figure online)

on the outside (inside) of the annulus. If so, the connection γ_u to γ_x must be connected through x_{in} , and γ_v to γ_y through v_{in} . However, as there is no cross-annulus connection between γ_u and γ_y , any connection from γ_u to x_{in} separates γ_y and v_{in} . Hence, both connections cannot exist at the same time. The red and blue disconnected components are thus on different sides of the annulus and we connect them to γ_u respectively γ_y without mutually interfering (see Fig. 7). \square

Corollary 1. *If a structure exists, then a structure also exists that has exactly one red and one blue connection across each annulus.*

Lemma 7. *If a structure exists, then there also exists a structure with exactly one red and one blue cross-annulus connection starting from any two regions on the inner annulus and connecting to any two regions on the outer annulus.*

Proof. Let an *interval* be a maximal arc of the same color on the boundary. By Corollary 1 we know there exists a structure with exactly one red and one blue connection across the annulus. Let γ_x be the blue cross-annulus connection and γ_y the red cross-annulus connection. We show that each of the endpoints of the cross-annulus connection can freely be moved. W.l.o.g., assume that γ_x is not counter-clockwise adjacent to γ_y on the outside of the annulus. Let k_{out} , l_{out} , and m_{out} be three intervals in clockwise order on the outer boundary of the annulus. We say a clockwise connection through the annulus from k_{out} to m_{out} covers l_{out} .

Let b_{out} be the blue interval that is counter-clockwise adjacent to y_{out} and r_{out} the red interval that is counter-clockwise adjacent to b_{out} . Interval b_{out} may have several incoming blue connections that cover r_{out} (see Fig. 8(a)). We can rewire the blue connections inside the annulus to connect the blue intervals in sorted order around the annulus, resulting in only one blue connection γ_b that covers r_{out} . Similarly we can also rewire the red connections covering r_{out} , and ending at y_{out} , to ensure only one red connection γ_r covers r_{out} .

Remove γ_y and insert a new red cross-annulus connection $\gamma_z = (y_{in}, r_{out})$. The connection γ_z can only intersect γ_r and γ_b . Removing γ_r results in two red components, one of which contains γ_z . Assume w.l.o.g. that y_{out} is in the same

connected component as γ_z . As γ_r intersected γ_z , the disconnected component can be connected to γ_z while only intersecting γ_b (see Fig. 8(b)).

Removing γ_b results in two blue components, one of which contains γ_x . We prove that b_{out} must be part of the blue component not containing γ_x . Assume to the contrary that b_{out} is still connected to γ_x . Interval r_{out} must be connected to y_{out} outside of the annulus as there was only one red cross-annulus connection and γ_b blocked any connection through the inside of the annulus. Similarly, interval b_{out} must have been connected through the outside of the annulus, as it is separated from any other region inside the annulus by γ_y and γ_z . However, they cannot both be connected through the outside of the annulus, as the connection r_{out} to y_{out} separates b_{out} and x_{out} on the outside of the annulus. Contradiction.

Therefore, we can safely reconnect the disconnected blue component through the annulus to γ_x (see Fig. 8(c)). Repeatedly moving the end-point of one of the cross-annulus connections allows the creation of any configuration of the two red and blue cross-annulus connections without invalidating the structure. \square

Lemma 7 implies that we can cut the annulus open to reduce the number of holes of a purple region by one without changing the solvability of the problem. Together with Lemma 4, this then implies Theorem 1.

4 Optimizing Panels

As shown, not all colored grids admit a painting. Here we investigate the design of the panels themselves, assuming that some painting is possible. To this end, we define the complexity of a panel as the number of *pieces* of maximal red and blue areas in the panel, see Fig. 9. The complexity of a painting is the maximal complexity of any of its panels. A t -panel (t -painting) has complexity t .




Fig. 9. Panels with complexity 3 and 5 respectively. (Color figure online)

Assuming some painting exists, we prove in this section that a 5-painting exists in general and that even a 2-painting exists if there are no white cells.

4.1 Ensuring a 5-Painting

We prove here that a 5-painting can always be realized. To this end, we show that a valid painting for a colored grid can be redrawn to include no more than three colored intervals along each side of all panels.

Lemma 8. *If a 2-colored grid admits a painting, then it admits a painting where each panel π has at most 3 intervals of alternating red and blue along each side.*

Proof (sketch). Assume that a panel π has at least 4 intervals of alternating red and blue on the left-side of π . As the painting is valid, both blue (/red) intervals are connected in the painting. For each interval we identify whether




Fig. 10. Reducing the number of intervals along a side of panel π , where there are at least four. (b) The two middle directions cannot be the same, as we cannot connect them with nonintersecting paths. (c) Shortcutting inside π reduces the number of intervals while maintaining a painting. (Color figure online)




Fig. 11. A panel with six pieces can always be reduced to have five, using either dotted line. (Color figure online)

the path exiting or entering π connects to the other interval of the same color (see Fig. 10(a)). The red and blue path cannot leave or exit the border of π in the same direction for the middle two intervals (see Fig. 10(b)). To reduce the number of intervals, we recolor the interval by shortcutting both the blue and the red piece inside π (see Fig. 10(c)). \square

Theorem 2. *If a partially 2-colored grid admits a painting, then it admits a 5-painting.*

Proof. By Lemma 8 there are at most three alternately colored intervals along each side of π . If a red and blue interval meet in a corner, we extend one in π around the corner to get four intervals and use Lemma 8 to reduce it back to at most three. If we have more than five pieces, a piece that has only one interval in π can be removed while maintaining a painting. Each remaining piece connects at least two intervals: with k intervals, the number of pieces is at most $\lfloor k/2 \rfloor$. A 6-panel thus requires 12 intervals: four equal-color (red) corners and a middle interval (blue) along each side. This enforces two pieces between the blue intervals, and one in each corner. However, we can now reduce the number of pieces to five, connecting either two blue pieces or two red corners (Fig. 11). \square

This bound is tight as a 5-panel may be required when the grid includes white cells (Fig. 12(a)). A 5-panel with at least two pieces of each color is never required—though such a 4-panel may be necessary (Fig. 12(b)). The above proof implies that there is only one option to create such a 5-panel: it has only two ways to connect the two blue pieces; both can be simplified to a 4-panel (Fig. 13).

4.2 Ensuring a 2-Painting

We show that a fully 2-colored grid (rectangular and without white cells) even admits a 2-painting, provided it admits any painting. As an intermediate step, we first prove that a painting exists that uses only one blue piece in any panel.

Lemma 9. *If a fully 2-colored grid admits a painting, then it admits a painting in which each panel has at most one blue piece.*




Fig. 12. Examples requiring complex panels. (a) A colored grid requiring a 5-panel. (b) A colored grid requiring a 4-panel with two pieces of both colors in the same cell. (Color figure online)




Fig. 13. There are two configurations for a 5-panel where both colors have at least two pieces. Both possible configuration can be simplified to a 4-panel. (Color figure online)

Proof (sketch). Since the grid admits a painting, we show how to modify the painting of each purple region P to ensure that the lemma holds. We first create a blue spanning forest in the panels of P that connects the panel-centers of adjacent panels. This ensures that each panel has exactly one blue piece inside, but may result in a disconnected blue polygon. However, since we know that a painting exists and the current solution maps to some forest in G_b , Lemma 1 implies that its dual G_r has a cycle around some tree in the forest. Hence, we can add connections between unconnected subpolygons to create a single blue polygon again, without disconnecting the red polygon. \square

The above construction relies on the alternation of the blue and red intervals along the boundary of P . As there are no white cells we can guarantee this alternating pattern. Indeed, the higher complexity with white cells is caused by long connections along a purple region’s boundary that are needed to achieve this alternating pattern for a partially colored grid (e.g., Fig. 12).

Theorem 3. *If a fully 2-colored grid admits a painting, then it admits a 2-painting.*

Proof (sketch). Since the fully 2-colored grid admits a painting, Lemma 9 implies that there is a painting Π where the panel for every purple cell contains only a single blue piece. For any purple cell with more than one red piece, we remove red pieces that only connect to one neighboring panel and recolor red corners




Fig. 14. Reducing panel complexity when there are two red corners along the same panel side. (a) The corners are connected via adjacent (or the same) sides of the panel: connect r_1 and r_2 , and recolor r_3 to blue. (b) The corners are connected via opposite sides: recolor r_1 to blue and connect r_3 and r_4 as well as b_1 and b_2 . (Color figure online)




Fig. 15. Two diagonally positioned red corners. The complexity of the panel can be reduced by introducing a red L -shape that connects all the red. (a) Reducing complexity if either p_1 or p_2 was blue. (b) Reducing complexity if both p_1 and p_2 were red. (Color figure online)

to blue if the other three cells incident to that corner have a red corner as well. Now, the pattern of the panel matches one of the following four cases.

1. There are two red corners r_1 and r_2 on the same side of the panel. The connecting path exits the current panel via the same side and enters either on the same or adjacent side (see Fig. 14(a)).
2. There are two red corners r_1 and r_2 on the same side of the panel. The connecting path exits the panel via opposite sides of the panel (see Fig. 14(b)). The blue piece connects only downwards in the panel below.
3. There are two red corners r_1 and r_2 that do not share a common side of the panel. In this case the other corners are blue, otherwise one of the two previous cases applies (see Fig. 15(a)). Furthermore, either p_1 or p_2 is blue.
4. There are two red corners r_1 and r_2 that do not share a common side of the panel (see Fig. 15(b)). Furthermore, both p_1 and p_2 are red.

We design a reduction rule for each case, as sketched in Figs. 14 and 15. Repeated application of the reduction rules, interlaced with the reduction of the number of red pieces in a panel, results in a 2-painting. \square

5 Conclusion

We took the first steps towards investigating a disjoint-polygons representation for visualizing set memberships (hypergraphs). We investigated the 2-color

version in which each element is positioned as a cell in a (unit-)grid. We showed how to test whether a disjoint-polygons representation is possible for a given 2-colored grid. Moreover, we proved that if such a representation is possible, then we can also bound the complexity of the corresponding “panels” (the coloring of a single cell). Each panel requires at most five colored pieces, and even only two pieces are sufficient when no white cells are present in the grid.

There are myriad options for further exploration. As not all grids admit a painting, we could study minimizing the number of polygons of the same color. We have not touched upon variants with more colors: does our approach readily generalize? However, considering the restrictions already in the studied 2-color variant, it seems likely that many practical instances do not admit a painting. If we allow rearranging elements, the 2-color variant becomes trivial, but is particularly interesting for multiple colors. Finally, we may consider the situation where some cells have no assigned set of colors but may be painted using any subset of the colors. Given enough such cells, the disjoint-polygons encoding can then represent more than Zykov-planar hypergraphs but cannot represent all planar supports.

Acknowledgments. The authors would like to thank Jason Dykes for fruitful discussions at an early stage of this research.

References

1. Akitaya, H.A., Löffler, M., Tóth, C.D.: Multi-colored spanning graphs. In: Hu, Y., Nöllenburg, M. (eds.) GD 2016. LNCS, vol. 9801, pp. 81–93. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-50106-2_7
2. Alper, B., Riche, N.H., Ramos, G., Czerwinski, M., Czerwinski, M.: Design study of LineSets, a novel set visualization technique. *IEEE Trans. Vis. Comput. Graph.* **17**(12), 2259–2267 (2011)
3. Alsallakh, B., Aigner, W., Miksch, S., Hauser, H.: Radial sets: interactive visual analysis of large overlapping sets. *IEEE Trans. Vis. Comput. Graph.* **19**(12), 2496–2505 (2013)
4. Alsallakh, B., Micallef, L., Aigner, W., Hauser, H., Miksch, S., Rodgers, P.: The state of the art of set visualization. *Comput. Graph. Forum* **35**(1), 234–260 (2016)
5. Baron, M.: A note on the historical development of logic diagrams: Leibniz, Euler and Venn. *Math. Gaz.* **53**(384), 113–125 (1969)
6. Buchin, K., van Kreveld, M., Meijer, H., Speckmann, B., Verbeek, K.: On planar supports for hypergraphs. *J. Graph Algorithms Appl.* **15**(4), 533–549 (2011)
7. Collins, C., Penn, G., Carpendale, S.: Bubble sets: revealing set relations with isocontours over existing visualizations. *IEEE Trans. Vis. Comput. Graph.* **15**(6), 1009–1016 (2009)
8. Dinkla, K., van Kreveld, M., Speckmann, B., Westenberg, M.: Kelp diagrams: point set membership visualization. *Comput. Graph. Forum* **31**(3pt1), 875–884 (2012)
9. Eppstein, D., Italiano, G., Tamassia, R., Tarjan, R., Westbrook, J., Yung, M.: Maintenance of a minimum spanning forest in a dynamic plane graph. *J. Algorithms* **13**(1), 33–54 (1992)
10. Riche, N.H., Dwyer, T.: Untangling Euler diagrams. *IEEE Trans. Vis. Comput. Graph.* **16**(6), 1090–1099 (2010)

11. Hurtado, F., Korman, M., van Kreveld, M., Löffler, M., Sacristán, V., Silveira, R.I., Speckmann, B.: Colored spanning graphs for set visualization. In: Wismath, S., Wolff, A. (eds.) GD 2013. LNCS, vol. 8242, pp. 280–291. Springer, Cham (2013). https://doi.org/10.1007/978-3-319-03841-4_25
12. Kaufmann, M., van Kreveld, M., Speckmann, B.: Subdivision drawings of hypergraphs. In: Tollis, I.G., Patrignani, M. (eds.) GD 2008. LNCS, vol. 5417, pp. 396–407. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-00219-9_39
13. Mäkinen, E.: How to draw a hypergraph. *Int. J. Comput. Math.* **34**, 177–185 (1990)
14. Meulemans, W., Dykes, J., Slingsby, A., Turkyay, C., Wood, J.: Small multiples with gaps. *IEEE Trans. Vis. Comput. Graph.* **23**(1), 381–390 (2017)
15. Meulemans, W., Riche, N.H., Speckmann, B., Alper, B., Dwyer, T., Dwyer, T.: KelpFusion: a hybrid set visualization technique. *IEEE Trans. Vis. Comput. Graph.* **19**(11), 1846–1858 (2013)
16. Purchase, H.: Metrics for graph drawing aesthetics. *J. Vis. Lang. Comput.* **13**(5), 501–516 (2002)
17. Purchase, H.C., Cohen, R.F., James, M.: Validating graph drawing aesthetics. In: Brandenburg, F.J. (ed.) GD 1995. LNCS, vol. 1027, pp. 435–446. Springer, Heidelberg (1996). <https://doi.org/10.1007/BFb0021827>
18. Simonetto, P., Auber, D.: Visualise undrawable Euler diagrams. In: Proceedings of the 12th Conference on Information Visualisation, pp. 594–599 (2008)
19. Simonetto, P., Auber, D., Archambault, D.: Fully automatic visualisation of overlapping sets. *Comput. Graph. Forum* **28**(3), 967–974 (2009)
20. Tutte, W.: *Graph Theory*. Addison-Wesley, Menlo Park (1984)
21. van Goethem, A., Kostitsyna, I., van Kreveld, M., Meulemans, W., Sondag, M., Wulms, J.: The painter's problem: covering a grid with colored connected polygons. Computing Research Repository, [arXiv:1709.00001](https://arxiv.org/abs/1709.00001) (2017)
22. Walsh, T.: Hypermaps versus bipartite maps. *J. Comb. Theor.* **18**, 155–163 (1975)
23. Wood, J., Dykes, J.: Spatially ordered treemaps. *IEEE Trans. Vis. Comput. Graph.* **14**(6), 1348–1355 (2008)